AD-A132 534 SOME REMARKS ON THE ROLE OF LOGIC IN KNOWLEDGE REPRESENTATION(U) BOLT BERANEK AND NEWMAN INC CAMBRIDGE MA D J ISRAEL JUL 83 BBN-5388 N00014-77-C-0371 1/1

UNCLASSIFIED F/G 5/7 NL

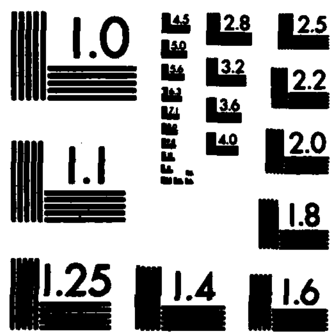

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD-A132 534

Bolt Beranek and Newman Inc.

12

Report No. 5388

Some Remarks on the Role of Logic in Knowledge Representation

Technical Report

David J. Israel

July 1983

Prepared for:
The Office of Naval Research

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER BBN Report No. 5388	2. GOVT ACCESSION NO. AD-A132 534	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) SOME REMARKS ON THE ROLE OF LOGIC IN KNOWLEDGE REPRESENTATION		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER BBN Report No. 5388
7. AUTHOR(s) David J. Israel		8. CONTRACT OR GRANT NUMBER(s) N00014-77-C-0371
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 10 Moulton St. Cambridge, MA 02238		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Department of the Navy Arlington, VA 22217		12. REPORT DATE July 1983
		13. NUMBER OF PAGES 20
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce, for sale to the general public.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Knowledge representation, logic, reasoning

The author presents

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

~~We attempt~~ a brief overview of some main aspects of the long-standing debate in Artificial Intelligence over the role of logic and logical formalisms in knowledge representation.

DD FORM 1473, 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Report No. 5388

SOME REMARKS ON THE ROLE OF LOGIC IN KNOWLEDGE REPRESENTATION

Technical Report

July 1983

David J. Israel

Prepared by:

Bolt Beranek and Newman Inc.
10 Moulton Street
Cambridge, Massachusetts 02238

Prepared for:

The Office of Naval Research

This research was supported by the Office of Naval Research under Contract No. N00014-77-C-0371. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Office of Naval Research or the U.S. Government.

TABLE OF CONTENTS

1. INTRODUCTION

If we date the founding of Artificial Intelligence as a field from the Dartmouth Conference of 1956, then it's fair to say that ever since its founding, there has been a debate among AI researchers about the appropriateness of adopting as a representational formalism a formal language of the kind devised, used and studied by logicians. And, if we are to believe SCIENCE magazine, the debate continues. Witness:

> Theoreticians... have reached no consensus on how to solve the AI problem - on how to make true thinking machines. Instead, there are two opposing philosophical viewpoints and a flurry of research activity along these two directions. The different viewpoints were represented at a recent meeting of the American Association for Artificial Intelligence by Marvin Minsky and John McCarthy...
>
> McCarthy believes that the way to solve the AI problem is to design computer programs to reason according to the well worked out languages of mathematical logic, whether or not that is actually the way people think. Minsky believes that a fruitful approach is to try to get computers to imitate the way the human mind works, which, he thinks, is almost certainly not with mathematical logic. [3]

The debate, clearly enough, is about the role that "logic" can play in solving "the AI problem". But not much beyond that is as clear. I will attempt in this paper to sort out <u>some</u> of the issues involved in this debate. In particular, I'm going to suggest that the failure to sort them out is one reason for the long life and inconclusive nature of the disagreement. First, however, something must be said by way of introduction.[1]

All parties to the debate agree that a central goal of research is that computers somehow or other come to "know" a good deal of what every human being knows about the world and about the organisms, natural or artificial, that inhabit it. This body of knowledge - indefinite, no doubt, in its boundaries - goes by the name of "common sense". How to impart such knowledge to a robot? Or: How to design a robot with a reasoning capacity sufficiently powerful and fruitful that when provided with some sub-body of this knowledge, it will be able to generate enough of the rest to enable the robot intelligently to adapt to and exploit its environment.[2]

I will assume that when parties to the debate speak of the "AI problem" or of the "knowledge representation problem", what they have in mind is the problem posed above. But what do they have in mind when they speak of "logic"?

---

[1]This essay is best seen as a minor addendum to [2] and [7] to both of which I am sorely indebted.

[2]It is assumed that most, if not all, common-sense knowledge is general, as in the knowledge that objects fall unless they are supported, that physical objects do not suddenly disappear, that one can get wet in the rain. It is, then, further assumed that knowledge of the particular facts of its current situation is the result of the deployment of the robot's sensors. The 3 examples of common sense knowledge are taken from [8].

2. WHAT'S LOGIC

What's logic? That's a question with no short answers. Suffice it say that what is at issue is the use of formal systems of the kind developed by mathematical logicians. These systems were originally created and studied with an eye toward the goal of a precisely characterizable symbolic language within which all mathematical propositions could be expressed. More particularly, a language within which one could express a set of basic mathematical truths or axioms from which all of the rest could be generated by the application of a finite set of precisely characterized (combinatorial) rules of proof which could be shown to be truth-preserving.

As just characterized, the languages of mathematical logic were not meant for "general" use. Their developers did not claim they were unrestrictedly universal symbolisms: that everything thinkable was adequately expressible in them. Indeed, it was not even held that everything sayable in a natural language was expressible in a formalized logical language. And surely much of common sense knowledge is expressible in (e.g.) English; witness the 3 examples drawn from Nilsson in the preceding footnote. Of course, that these formalisms were not devised with an eye toward the expression of common sense knowledge, that is, toward solving the AI problem, doesn't signify that they can't be so used. To argue this last point, one must demonstrate particular failures of representational adequacy, perhaps with explanations of said failures in terms of the differences between mathematical knowledge and common sense knowledge.

<u>Some</u> attempts have been made in this direction. Doubts have been raised about the adequacy or appropriateness, for example,

of the "basic" language of logic - the language of the first-order predicate calculus. Of course, it is not obvious that these translate into doubts about the adequacy of logical languages in general. There is nothing in the spirit of the McCarthy camp to rule out the use of many different logical languages in addressing the representation problem. Moreover many of these doubts have been based on criticisms of some particular way of formally representing some body of knowledge. Such considerations are best understood as objections to the effect that the objects, properties and relations of the domain have been wrongly conceived or wrongly represented. Instead, they seem to have been understood as objections to the language of logic itself; as if adopting some logical language involved adopting some one particular way of "cutting the world at its joints", a way that was arguably inappropriate to many domains and for many purposes. As a representational formalism, a logical language is (just) a tool. Given a certain task, this tool can be used more or less well. A commitment to standard logical formalisms does not carry along with it a commitment to a particular metaphysics or ontology, let alone to a particularly wrong-headed one.

For the most part, those in the Minsky camp seem to assume that the presumed enormity of the differences between mathematics and common sense, together with the fact that the formalisms of mathematical logic were meant for and are fine for the former, somehow guarantees the inadequacy of such languages for the latter. Moreover this guarantee is taken to absolve "the enemies of logic" of the responsibility actually to show this principled inadequacy in any detail.

Before surveying, though briefly, some of the alleged "inherited" maladaptive traits of logical formalisms, it will be well to separate out at least three different aspects of such formalisms. In the first place we have a formal language. This is determined by specifying a vocabulary, broken down into syntactic types, and a set of formation rules for generating complex expressions, in particular, complex well-formed terms and/or formulae of the language. All of these specifications are to be mechanical or algorithmic.

Having determined the syntax of a language, one can then go on to specify a semantics for the language in a way that mirrors the recursive specification of its syntax. Thus one specifies an assignment of semantic values to the primitive non-logical expressions of the language, different types of semantic value being associated with the different syntactic types. Such assignments are also called models or interpretations. One then gives a set of rules of interpretation which determine semantic values of complex expressions as a function of the semantic values of their constituents and the syntactic formation rules used to generate them. For instance, the meaning of the logical constants - e.g., the truth-functional connectives and quantifiers - can be given by rules of interpretation which assign values to complex expressions containing them.

Finally, we may specify a deductive apparatus for the language. This can be done in many different ways, but at its heart is the specification of a set of rules of transformation, whose applicability to a set of sentences can be effectively determined and whose output, typically a sentence, is likewise determinable. Moreover, these rules must "jibe" with the semantics we have specified for our language. For example: if the premisses of a rule are valid - true in every interpretation, according to our semantics - then so too is the conclusion of the rule. This is the requirement that the rules be sound with respect to validity. We might also require that the rules be truth-preserving.

3. ON BEING LOGICAL

Now that we've gone to the trouble of characterizing a logical system, what more can be said about the issues dividing McCarthy and Minsky? A crucial point to make (and already alluded to above) is that certain historically very important objections to the applicability of logic to the AI problem have not really been aimed at the expressive capabilities of logical languages. Rather they have been directed at the claim - to which the "friends of logic" are supposed to be committed - that <u>common sense reasoning</u> or <u>inference</u> could be adequately captured by running a sound theorem-prover over such a language. It is on these that I intend to concentrate. Indeed much of [3] is devoted to a particular instance of this mode of objection - the problem of default reasoning or of "non-monotonic logic".

Remember the claim:" ...the way the human mind works which is almost certainly not with mathematical logic." Throughout [3] there is much talk, on both sides, of mathematical and common-sense <u>reasoning</u>. Thus, McCarthy answers Minsky:

> "Minsky never liked logic," says McCarthy. "when difficulties with mathematical reasoning came up, he felt they killed off logic. Those of us who did like logic thought we should find a way of fixing the difficulties." Whether logical reasoning is really the way the brain works is beside the point, McCarthy says. "This is A(RTIFICIAL) I(ntelligence) and so we don't care if it's psychologically real".

The same tendency is evident in Minsky's recent AI Magazine paper "Why People Think Computers Can't" [5].

> Many AI workers have continued to pursue the use of logic to solve problems. This hasn't worked very well,

> in my opinion; logical reasoning is more appropriate for displaying or confirming the <u>results</u> of thinking than for thinking itself. That is, I suspect we use it less for solving problems than we use it for explaining the solutions to other people and - much more important - to ourselves.

It is as if it were assumed on <u>both sides</u> that the commitment to use a logical language as a <u>representation</u> language carried with it ineluctably a commitment to some sound algorithmic deductive apparatus as the central or sole non-perceptual generator of new knowledge or new beliefs. But the first is quite independent of the second; the second, much more contentious than the first. Having specified a formal, logical language and its semantics, one is free to specify any rules of transformation one likes or believes useful. They need not be sound; they need only be mechanically applicable. That is, the conditions of their "legal" applicability must be decidable solely in virtue of the syntactic structures of sentences. For instance, we might be able to come up with such rules which embody useful principles of plausible or probabilistic reasoning, or even of analogical reasoning. Again, just as they need not be sound, they need not be non-domain-specific. <u>The applicability of the rules can depend on the occurrences of particular non-logical, descriptive expressions in sentences</u> as well as on the occurrences of the logical constants.

Holding that such rules need not be sound does not commit one to ignoring the semantics of the language in one's specification of the inferential procedures that operate on its sentences. It simply frees one to experiment. What we want is a set of rules which, collectively, embody - in syntactically codifiable form - fruitful and generally reliable modes of reasoning. They need not be our own; and they certainly need not be discovered by introspection. They must, however, on reflection seem reasonable to us. But it's hard to see how to go about devising mechanizable rules that embody rational principles of belief fixation and revision if one doesn't have a firm grasp of the meanings of the sentences on which those rules act. This last consideration immediately suggests one very strong argument for using a formal logical language: we can get our hands on precise accounts of what their sentences mean.

4. REASONING AND LOGIC

To bring out the freedom we have in our choice of rules, we must be clear about the distinction between reasoning and proof. Let's begin by homing in on a "psychological" issue. Minsky seems to be claiming that logic - by which he means the application of deductively sound rules of proof - doesn't play an important part in common sense reasoning. McCarthy is prepared to be neutral on this point. So far, so good; but surely this suggests that the parties might agree that a distinction can be made between reasoning and proof. And they're right. Exemplary reasoning can often lead us from true beliefs to false ones. Reasoning often involves going out on a limb a little, going beyond what we are absolutely sure of or take for granted. Indeed, reasoning can often lead us to give up some of the beliefs from which we began; this can happen even when we have not set out purposefully to put those beliefs to the test. (Contrast this with proofs by refutation or with reductio ad absurdum proofs in logic.)

To take a simple case of modus ponens: suppose you accept - among other things, of course - some sentence of the form if P then Q and accept the antecedent. Should you, need you, accept the consequent? Not necessarily, for you have may tremendously good <u>over-all</u> reasons for believing not-Q and these might lead you to give up belief in either the conditional or its antecedent. Further, rules of proof are local; they apply to a given set of sentences in virtue of their individual syntactic forms. Reasoning, on the other hand, can often be global - one must try to take into account all the relevant evidence in one's possession. One must often try to get more evidence if that on hand is judged insufficient. That judgment, and judgments about the relevance and weights of evidence, are typically themselves the products of reasoning.

It may appear that logical proof is being opposed to reasoning. The correct view seems to be that logical proof is a tool used in reasoning. Hence, the inappropriateness of talk about "logical reasoning", especially if it carries the connotation that the real reasoning that proof subserves is "illogical" or non-logical.

Two further points should be made. If one looks at the history of science, one can see part of the force of Minsky's claim that "logical reasoning [sic] is more appropriate for displaying or confirming the <u>results</u> of thinking than for thinking itself". Axiomatic formalizations by logicians of a body of knowledge - when such formalizations are forthcoming at all - come after the scientists have done their work. But this does not mean that proof, deductive inference, has played no part, or only an insignificant part, in the scientific work. Moreover, with respect to the particular case at issue, that of representing common-sense knowledge, the "science" has long since done most of its job; most of us already know a whole lot about the way the world works. It's not that there are lots of new discoveries to be made; although it's also surely not that there are none. A significant part of the problem is precisely to codify and systematize this knowledge. There <u>may</u> be reasons for doubting the adequacy of logical formalisms for <u>this</u> task; but these need have nothing to do with the claims about the adequacy of deductively sound rules of proof.

Note, though, that the claim that deductively valid rules of proof are all that is required is an extraordinarily strong one. It imposes an extremely strong condition of adequacy on a formalization of common-sense knowledge; namely, that everything a robot needs to know, even in a constrained, but real environment, is a deductive consequence of the things we "tell" it - together with the particular facts delivered by its sensors. For on such an account, the only way a robot can learn new things, except those it learns by perceiving, is by deducing them from what it already knows. Quite independent of the question of how people do it, it's going to be awfully hard for us to arrange it so that that's how successful artificially intelligent beings do it. This, of course, is not an argument that we should give up trying to formulate as much as we can as systematically as we can.

5. "NON-MONOTONIC LOGIC"

Much of [3] is devoted to a discussion of one(?) particular species of common sense reasoning that is supposed to be beyond the purview of logic. This is described as the problem of dealing with exceptions. Minsky sees this as one of those defects of "logic" that can be traced to its origins - and an especially glaring defect, it is.

> Logical systems work very well in mathematics, but that is a well-defined world. The only time when you can say something like, If a and b are integers, then a plus b always equals b plus a, is in mathematics... Consider a fact like, 'Birds can fly.' If you think that common-sense reasoning is like logical reasoning then you believe there are general principles that state, 'If Joe is a bird and birds can fly, then Joe can fly.' Suppose Joe is an ostrich or a penguin? Well, we can axiomatize and say if Joe is a bird and Joe is not an ostrich or a penguin, then Joe can fly. But suppose Joe is dead? Or suppose Joe has his feet set in concrete? The problem with logic is once you deduce something you can't get rid of it. What I'm getting at is that there is a problem with exceptions. It is very hard to find things that are always true.

This problem with logic is alleged to be due to monotonicity - if a sentence S is a logical consequence of a set of sentences A, then S is (still) a logical consequence of any set of sentences that includes A. So, if one thinks of A as embodying the set of beliefs with which one starts, the addition of new beliefs cannot lead to the "logical" repudiation of old consequences. ("Once you deduce something, you can't get rid of it.") Thus McCarthy:

> A proper axiomatization is one in which a proof exists for all conclusions that are ordinarily drawn from these facts. But what we know now about common sense is

> that that's asking for too much. You need another kind of reasoning - nonmonotonic reasoning...If you know I have a car, you may conclude that you can ask me for a ride. If I tell you the car is in the shop, you may conclude you can't ask me for a ride. If I tell you it will be out of the shop in 2 hrs., you may conclude you can ask me. [As more premises are added, the conclusion keeps changing.]

By my lights, the alleged defect of logic is no defect at all; indeed, it has nothing directly to do with logic. Logic doesn't tell you what to hold on to; nor what to get rid of. That's the job of reasoning - which is surely a non-monotonic process. Finding out or coming to believe new things often gives us good reason for repudiating old favorites; remember the case of modus ponens. Throwing away one's initial beliefs is in no way illogical - especially not if you throw them away because what they logically entail conflicts with what you have overwhelming reasons to believe. This is one reason it's misleading to talk of "premises" in reasoning. Another reason has to do with the global nature of reasoning - the fact that, in principle, it's nothing less than a whole theory that operates as what's given, and, in practice, often good size chunks of theory. Surely it's odd to think of a theory - especially construed as closed under consequence - as a premise.

Just as logic doesn't tell us what beliefs to keep, it doesn't tell us what beliefs to throw away. In particular, it doesn't tell us what to do when we discover, with its help, that we hold inconsistent beliefs. It only tells us that in that case not all of our beliefs can be true. The fact that, in many standard logics, from a contradiction anything and everything follows is quite irrelevant - so long, that is, as one distinguishes clearly between logic and reasoning, and sees the former as a tool used in the latter. (There are, moreover, logics defined over perfectly standard languages in which it is not the case that everything follows from a contradiction.)

The problem to which Minsky and McCarthy are addressing themselves is a deep one and McCarthy, in particular,[3] has made a significant technical contribution in the area [4]. But the crucial point is that nothing in the debate about non-monotonic logic argues against the use of a standard logical language, with some standard semantic account, as a representation language for Artificial Intelligence. So long, that is, as "logic" is kept in its proper place. As Minsky himself says:

> But "Logic" simply isn't a theory of reasoning at all. It doesn't even try to describe how a reasoning process might work. It is only a partial theory of how to <u>constrain</u> such a process... [6]

If it doesn't even try, it can scarcely be said to botch the job, can it?

---

[3]but not alone: see [1]

6. CONCLUSION

> All efforts to solve the knowledge representation problem share two major obstacles, McCarthy explains. "The preliminary problem is to decide what knowledge to represent. The key thing that we have not got formulated [are] the facts of the common sense world." Then, even if researchers do manage to represent knowledge (common sense - D.I) in computers, they still are faced with the problem of getting answers out of the computer in a reasonable time. [3].

Though I have limited myself to discussing only a few facets of the disagreement between McCarthy and Minsky, I should like to second McCarthy's characterization of the representation problem - as I'm sure Minsky would do as well. Before we worry over much about the adequacy of one or another representational formalism, we should have some better ideas about what it is we want to represent. My intuition is that the more self-consciously and systematically we set about the task of figuring out what we all know about the world and its ways, the more fully we make explicit our usually implicit background beliefs - the beliefs we usually take for granted; the beliefs too obvious even to mention - the more likely will it seem that the problem is not one of the expressive power of this or that formalism. Rather, the problem will be that of figuring out reasonable ways of applying the relevant parts of what we know to a given problem - and this, in a world in which, though the same thing never happens twice, success depends on grasping whatever reliable and exploitable regularities there are.

REFERENCES

[1]
Special Issue on Non-Monotonic Logic.
Artificial Intelligence Vol. 13, nos. 1,2, April 1980.

[2] Hayes, P. J.
In Defence of Logic.
In *Proc. IJCAI-77*, pages 559-565. International Joint Conferences on Artificial Intelligence, Cambridge, MA, 1977.

[3] Kolata, G.
How Can Computers Get Common Sense.
Science 217:1237-1238, September 24, 1982.

[4] McCarthy, J.
Circumscription - A Form of Non-Monotonic Reasoning.
Artificial Intelligence 13(1,2):27-39, April, 1980.

[5] Minsky, M.
Why People Think Computers Can't.
AI Magazine III(4):3-15, Fall, 1982.

[6] Minsky, M.
Learning Meaning.
unpublished manuscript.

[7] Moore, R. C.
The Role of Logic in Knowledge Representation and Commonsense Reasoning.
In *Proceedings of the National Conference on Artificial Intelligence*, pages 428-433. American Association for Artificial Intelligence, August, 1982.

[8] Nilsson, N.
Artificial Intelligence: Engineering, Science, or Slogan?
AI Magazine 3(1):2-8, 1982.

Official Distribution List

Contract N00014-77-C-0371

	Copies
Defense Documentation Center Cameron Station Alexandria, VA 22314	12
Office of Naval Research Information Systems Program Code 437 Arlington, VA 22217	2
Office of Naval Research Code 200 Arlington, VA 22217	1
Office of Naval Research Code 455 Arlington, VA 22217	1
Office of Naval Research Code 458 Arlington, VA 22217	1
Office of Naval Research Branch Office, Boston 495 Summer Street Boston, MA 02210	1
Office of Naval Research Branch Office, Chicago 536 South Clark Street Chicago, IL 60605	1
Office of Naval Research Branch Office, Pasadena 1030 East Green Street Pasadena, CA 91106	1
Naval Research Laboratory Technical Information Division Code 2627 Washington, D.C. 20380	6

cont'd.

Naval Ocean Systems Center 1
Advanced Software Technology Division
Code 5200
San Diego, CA 92152

Dr. A. L. Slafkosky 1
Scientific Advisor
Commandant of the Marine Corps
(Code RD-1)
Washington, D.C. 20380

Mr. E. H. Gleissner 1
Naval Ship Research & Development Ctr.
Computation & Mathematics Dept.
Bethesda, MD 20084

Capt. Grace M. Hopper, USNR 1
Naval Data Automation Command
Code 00H
Washington Navy Yard
Washington, D.C. 20374

Mr. Paul M. Robinson, Jr. 1
NAVDAC 33
Washington Navy Yard
Washington, D.C. 20374

Advanced Research Projects Agency 1
Information Processing Techniques
1400 Wilson Boulevard
Arlington, VA 22209

Capt. Richard L. Martin, USN 1
507 Breezy Point Crescent
Norfolk, VA 23511

Director, National Security Agency 1
Attn: R54, Mr. Page
Fort G.G. Meade, MD 20755

Director, National Security Agency 1
Attn: R54, Mr. Glick
Fort G.G. Meade, MD 20755

Major James R. Kreer 1
Chief, Information Sciences
Dept. of the Air Force
Air Force Office of Scientific Research
European Office of Aerospace
Research & Development
Box 14
FPO New York 09510

cont'd.

Mr. Fred M. Griffee
Technical Advisor C3 Division
Marine Corps Development
& Education Command
Quantico, VA 22134 1

END

FILMED

9-83

DTIC